arXiv:1408.2812v1 [cs.CC] 12 Aug 2014

Homomorphism reconfiguration via homotopy

Marcin Wrochna

We consider the following problem for a fixed graph H: given a graph G and two H-colorings of G, i.e. homomorphisms from G to H, can one be transformed (reconfigured) into the other by changing one color at a time, maintaining an H-coloring throughout. This is the same as finding a path in the $\mathrm{Hom}(G, H)$ complex. For $H = K_k$ this is the problem of finding paths between k-colorings, which was shown to be in P for $k \leq 3$ and PSPACE-complete otherwise by Cereceda et al. 2011. We generalize the positive side of this dichotomy by providing an algorithm that solves the problem in polynomial time for any H with no C_4 subgraph. This gives a large class of constraints for which finding solutions to the Constraint Satisfaction Problem is NP-complete, but finding paths in the solution space is P.

The algorithm uses a characterization of possible reconfiguration sequences (paths in $\mathrm{Hom}(G, H)$), whose main part is a purely topological condition described in algebraic terms of the fundamental groupoid of H seen as a topological space.

1. Introduction

Reconfiguration

Reconfiguration is a framework in which we study how discrete structures, constrained in various ways, can be carefully transformed with small steps. This is often best described by finding paths in a solution graph, whose vertices are all solutions to a combinatorial problem and whose edges define the steps between solutions one is allowed to make.

For example, in k-Recoloring [BC09; CHJ11; Joh+14; BM14], one is given two proper k-colorings of a graph G and the question is whether one can be transformed into another by changing one color at a time, maintaining a proper coloring throughout. In other words, the solutions are proper k-colorings and a reconfiguration step is changing the color of one vertex. Another well studied example is Token Jumping [HD05; KMM12; BKW14], where the solutions are independent sets of some given size (seen as sets of tokens on the graph's vertices) and a reconfiguration step removes one vertex from the set to add another (jumps one token to a different vertex). Yet another example is the reconfiguration of generalized SAT problems [Gop+09; Sch13; Mou+14], where solutions are satisfying assignments of a given formula, and a reconfiguration step flips one variable of the assignment.

Homomorphisms

Our aim is to give more general statements about reconfiguration by defining a problem parameterized by the different constraints allowed. The natural tool for expressing systems being reconfigured seems to be Constraint Satisfaction Problems. This generalizes k-Coloring and generalized SAT problems, while (weighted, counting, ...) variants generalize Independent Set, for example. However, graph homomorphisms form an elegant subclass of problems, which is known to express many features of general CSPs. This is why we consider homomorphism reconfiguration. Indeed, this approach allowed us to argue in [Wro14b] that the only notion of sparseness that can apply to (unparameterized) reconfiguration problems in general is treedepth, and that many of them remain PSPACE-complete even in graphs of bounded bandwidth (thus pathwidth, treewidth, etc.). The reduction showed there explains why reconfiguration variants of easy combinatorial problems can be hard. In this paper we explore why reconfiguration variants of hard problems can be easy.

A homomorphism from a graph G to a graph H is a mapping $\sigma : V(G) \to V(H)$ such that edges are mapped to edges, i.e., $uv \in E(G)$ implies $\sigma(u)\sigma(v) \in E(H)$. We also use the name H-coloring, especially when H is fixed in the context. Vertices of H are then called *colors*. Note that a K_k-coloring (so H is the graph with all edges except loops) is the same as a k-coloring.

Define the solution graph $\mathrm{Hom}_1(G, H)$ to be the graph with H-colorings of G as vertices and edges between any two H-colorings that differ in the color of only one vertex (the name comes from it being the 1-skeleton of the $\mathrm{Hom}(G, H)$ complex, see below). For a fixed graph H, H-Recoloring is the problem asking whether two given H-colorings of a given graph G are connected by a path in $\mathrm{Hom}_1(G, H)$. Shortest H-Recoloring asks whether there is such a path of at most some given length.

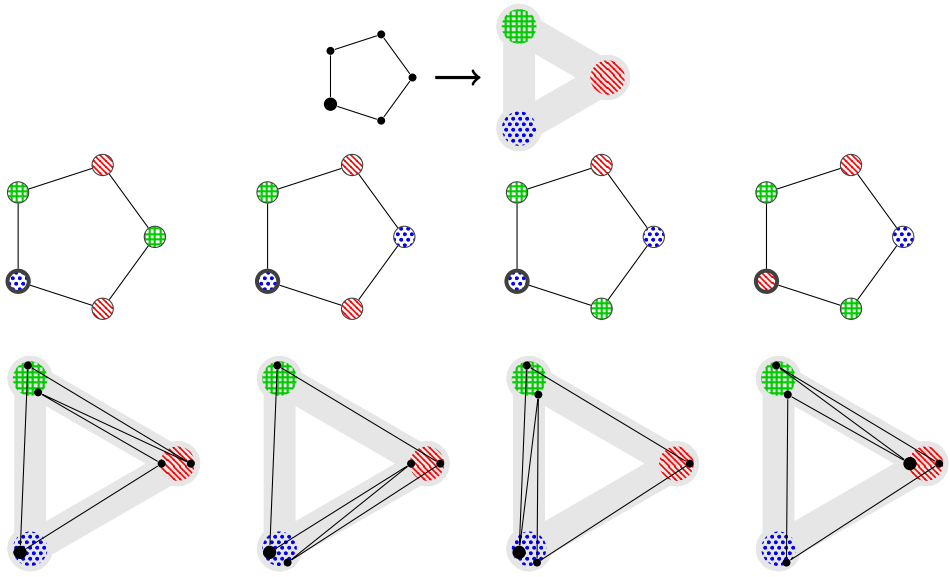

Figure 1: A sequence of 3-colorings of C_5 and the same sequence seen as a K_3-recoloring sequence of homomorphisms from C_5 to K_3 (a graph with three vertices: red, green, blue). One vertex of C_5 is thickened for clarity.

Motivations

The primary motivation for studying reconfiguration problems is to investigate the solution space of combinatorial problems, especially from the perspective of local search heuristics and random solution sampling. In particular, the success of Survey Propagation as a method for solving random Constraint Satisfaction Problems (CSPs) is connected to several conjectures about the structure of clusters of satisfying assignments (connected components in the solution graph) and frozen variables (variables/vertices that cannot change their value/color by any sequence of steps), see [ACORT11] (a reference the author would recommend to anyone interested in reconfiguration).

While finding paths in the solution graph is, for the above purposes, mostly a toy problem, it arises more directly in some settings, such as puzzles with sliding blocks. Indeed, the Nondeterministic Constraint Logic construction of [HD05] allowed to show that many popular puzzles are PSPACE-complete [HD09], and more interestingly, so is the equivalence of proofs in a certain proof system [HH14], answering a question about normal forms of proofs that arose in this context.

Related work

H-Recoloring is known to be PSPACE-complete (as a decision problem) for $H = K_k$ where $k \geq 4$ [BC09] and in P for $k \leq 3$ [CHJ11]. The latter result has been improved to show that Shortest K_3-Recoloring is also in P [Joh+14].

For CSPs in the boolean domain, a dichotomy was shown by [Gop+09] – for a fixed set of boolean constraints Γ (boolean relations, clause types), the problem of finding paths in the solution graph of a SAT(Γ) instance is either in P or PSPACE-complete. In particular it is always in P when the corresponding satisfiability problem is in P (e.g. 2-SAT or Horn-SAT), but it is also in P for some Γ for which satisfiability is NP-complete. This was slightly corrected and extended to several similar problems by Schwerdtfeger [Sch13], while a trichotomy was shown for the problem of finding *shortest* paths by [Mou+14]. Both [Gop+09] and [Sch13] asked whether their results could be extended to larger domains. Our work can be seen as a step in this direction, but limited to only one symmetric relation of arity 2.

The corresponding extension of Schaefer's dichotomy [Sch78] for satisfiability (deciding the existence of a solution) to CSPs with arbitrary finite domains is a long-standing open problem stated by Feder and Vardi [FV98]. They showed that the conjecture is equivalent when limited to one relation of arity 2 (digraph homomorphisms). Hell and Nešetřil proved the dichotomy when the relation is additionally assumed to be symmetic (graph homomorphism) [HN90]: the problem of deciding the existence of an H-coloring of a given graph is in P for H bipartite or containing a loop, and NP-complete otherwise.

Results

It is natural to ask whether the unexpected tractability of K_3-Recoloring (in light of 3-Colorability being NP-complete, even in 4-regular planar graphs [Dai80]) is caused by the following property: whenever a vertex changes its color in a 3-coloring (e.g. from red to green), all of its neighbors must have one common color (blue).

We answer this in the positive considering the following definition: a graph H has the *monochromatic neighborhood property* if for every two colors $a, b \in V(H)$, the set of common neighbors $N_H(a) \cap N_H(b)$ contains at most one color. For graphs without loops this is equivalent to not containing C_4 (the cycle on 4 vertices) as a subgraph (not necessarily induced). This includes K_3 and all graphs of girth ≥ 5, for example. For graphs with loops allowed, this is equivalent to not containing C_4, K_3 with one loop added, nor K_2 with both loops added.

We show an algorithm that solves H-Recoloring in polynomial time for all H with the monochromatic neighborhood property. To achieve this, we characterize possible paths in $\mathrm{Hom}_1(G, H)$ by describing sequences of colors which can occur when we observe one given vertex of G throughout an H-recoloring. It turns out they are mostly limited by the fact that reconfiguring an H-recoloring corresponds to a continuous transformation of a map from G to H as topological spaces.

In combinatorial algebraic topology

Reconfiguration of homomorphisms has already been studied independently in the field of combinatorial algebraic topology, though from a different angle. The notion of $\times$-homotopy of homomorphisms as defined by Dochtermann [Doc09] is identical to reachability in the solution graph. The solution graph of homomorphisms from G to H arises as the subgraph of the exponential graph H^G induced by looped vertices and as the 1-skeleton of the Hom-complex $\mathrm{Hom}(G, H)$, see [BK06]. The Hom-complex, studied for its interesting categorical and topological properties, is a construction similar (and homotopy equivalent) to the clique complex of our solution graph. These definitions were introduced to provide lower bounds on the chromatic number of graphs, a notoriously hard problem. A typical theorem derived from such methods is that for loopless graphs G, H, if the solution graph of $G \to H$ homomorphisms is connected for all G of degree at most d, then the chromatic number of H is at least $d/2$ (and is conjectured to be at least d) [BW04]. Studies have thus been mostly concerned with highly regular graphs for which the solution graph can be proved to be in some sense tightly connected.

We do not know of any prior work on the computational complexity of deciding $\times$-homotopy. More interestingly, we do not know of any prior example where $\times$-homotopy would be related to the continuous transformation of continuous mappings between graphs, more than through (functorial) analogy.

2. Preliminaries

An (undirected) graph G is a pair $(V(G), E(G))$ where $V(G)$ is a finite set of *vertices*, while $E(G)$ is the set of *edges* – undirected vertex pairs $\{u, v\}$, $u, v \in V(G)$, written uv for short. A *loop* is an edge uu for some vertex u. The neighborhood $N_G(v)$ is defined as $\{w \in V(G) \mid vw \in E(G)\}$. Hence $a \in N_H(a)$ iff H has a loop at a.

G and H in this paper are always connected undirected graphs with at least one edge. G is always assumed to have no loops. H can have loops, but is always assumed to have the monochromatic neighborhood property.

An *H-recoloring sequence* or *reconfiguration sequence* is a path in $\mathrm{Hom}_1(G, H)$, that is, a sequence of H-colorings of G in which consecutive colorings differ at one vertex. Consider a step of an H-recoloring sequence – a vertex $v \in V(G)$ changes color from $a \in V(H)$ to $b \in V(H)$. Since G is connected, loopless and has an edge, v has a neighbor, say $w \neq v$. As only v changes its color in the step, w has the same color, say $h \in V(H)$, before and after the step. The H-coloring before the step implies that $ha \in E(H)$, while the one after the step implies that $hb \in E(H)$. Thus $h \in N_H(a) \cap N_H(b)$. From the monochromatic neighborhood property we imply that $N_H(a) \cap N_H(b) = \{h\}$. We will often call h '*the color that all neighbors of v have during the step*' (that is, in the H-colorings before and after the step), without arguing its existence and uniqueness anymore.

3. Fundamental groupoid

An *oriented edge* of a graph H is an oriented pair $e = (u, v)$ such that $\{u, v\}$ is an edge of H; we denote its initial vertex u as $\iota(e)$ and its target vertex v as $\tau(e)$. We write e^{-1} for $(\tau(e), \iota(e))$. A *walk* from u to v in a graph H is a sequence of oriented edges $e_1, \ldots, e_l$ of H such that $\iota(e_1) = u$, $\tau(e_l) = v$ and $\tau(e_i) = \iota(e_{i+1})$ for $i = 0, \ldots, l-1$. We write ε for an empty walk (though some define a different empty walk for every vertex, the endpoints of ε will be irrelevant for us). The length of a walk is the number of edges in it. A walk W_1 from u to v can be concatenated to a walk W_2 from v to w to form a walk W_1W_2 from u to w.

We call a walk *reduced* if it contains no two consecutive edges e_ie_{i+1} such that $e_{i+1} = e_i^{-1}$. One can *reduce* a walk by removing any such two consecutive edges from the sequence. It can easily be seen that by iteratively reducing a walk W, one always gets the same reduced walk, which we therefore denote as $\overline{W}$. For any two reduced walks $\overline{W_1}, \overline{W_2}$ such that W_2 starts where W_1 ends, we write $\overline{W_1} \cdot \overline{W_2}$ for $\overline{\overline{W_1}\,\overline{W_2}}$ and similarly one can observe that $\cdot$ is associative. For any walk $W = e_1e_2 \ldots e_l$ we write W^{-1} for the reversed walk $e_l^{-1} \ldots e_2^{-1}e_1^{-1}$. Of course $\overline{W} \cdot \overline{W}^{-1} = \varepsilon = \overline{W}^{-1} \cdot \overline{W}$ and $\varepsilon \cdot \overline{W} = \overline{W} \cdot \varepsilon = \overline{W}$, so the set of non-backtracking walks of a graph forms together with the operations $\cdot$ and $()^{-1}$ what we call a *groupoid*. This particular groupoid is called the *fundamental groupoid* $\pi(H)$ of H.

Note that if H has a loop at u, then $(u, u) = (u, u)^{-1}$ is an oriented edge and a reduced walk of length 1. On the other hand, $(u, u) \cdot (u, u) = \varepsilon$. (Hence $\pi(H)$ is strictly speaking not a *free groupoid* and loops have no topological interpretation, see [KN07]).

A groupoid is in many ways similar to a group, except that the group operation $\cdot$ is a partial function, defined only when the 'head' of one element matches the 'tail' of the other.

More aptly, a groupoid could be defined as *a group with many objects*, that is, a category in which every morphism is invertible. Groupoids behave similarly to groups (much more so than semigroups) and identities such as $(e \cdot f)^{-1} = f^{-1} \cdot e^{-1}$ known from group theory are easily reproved in groupoids. Using the groupoid as opposed to the better known group allows us to describe calculations much more uniformly, without the tedious change of base points, for example.

Topological interpretation

Let us comment on how this algebraic structure captures the topology of curves in the graph. When referring to topology, continuous maps, curves and homotopy, we give informal interpretations without proof. Here we recall some classical results and definitions, which we do not require formally, but are helpful, if not crucial, in understanding our results.

A graph H without loops can be naturally associated with a topological space, constructed from copies of $[0, 1]$ for each edge, with endpoints merged into vertices accordingly. A *curve* in this space is a continuous map $f : [0, 1] \to H$ (where H is meant as a topological space), not necessarily injective. Two curves f_0, f_1 are *homotopic* if one can be continuously transformed into the other, which means there is a set of functions $\{\phi_t : t \in [0, 1]\}$ such that $\phi_0 = f_0, \phi_1 = f_1$ and the mapping $\phi : (t, x) \mapsto \phi_t(x)$ is continuous as a function from $[0, 1] \times [0, 1]$ to H.

The fundamental groupoid fully describes curves up to homotopy. For any two vertices u, v of H, two curves f_0, f_1 are homotopic via a homotopy ϕ_t that fixes the endpoints ($\phi_t(0) = u, \phi_t(1) = v$ for all t) if and only if the corresponding reduced walks in $\pi(H)$ are identical. Considering only curves (or equivalently, walks) that start and end in the same vertex v, we obtain a group, which is known as the fundamental group $\pi_1(H, v)$ of the topological space H. When no vertex is fixed, a closed curve starting and ending in v_1 is homotopic to a closed curve starting and ending in v_2 (via a homotopy such that $\phi_t(1) = \phi_t(0)$ for all t) if and only if the corresponding elements C_1, C_2 of $\pi(H)$ are *conjugate*, i.e. $C_2 = P^{-1} \cdot C_1 \cdot P$ for some $P \in \pi(H)$.

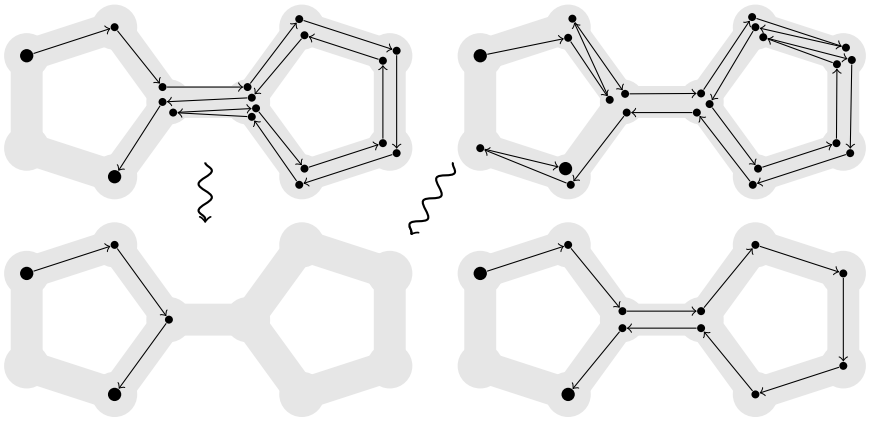

Figure 2: Examples of two walks (in a graph H on 10 vertices) which reduce to the same, bottom left one. The bottom right one is a different reduced walk; when its endpoints are fixed, it cannot be distorted as a curve to give any of the others.

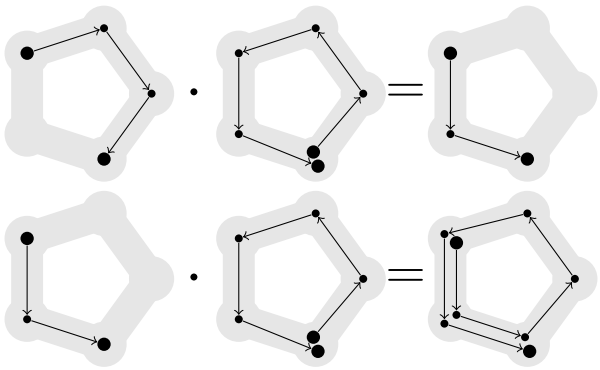

Figure 3: Examples of $\cdot$ multiplication in the fundamental groupoid of $H = C_5$.

4. Vertex walks and realizability

As a vertex $v \in V(G)$ changes colors from a to b in a step of an H-recoloring process, if h is the color that all neighbors of v have during the change, then $(a, h)(h, b)$ is a walk in H (of length 2). Looking at all the color changes of one vertex this way gives a walk in H which traces the colors that v had. This walk (for one arbitrary vertex v), even after reducing, will be shown to almost completely describe the H-recoloring sequence.

Formally, consider an H-recoloring sequence $S = \sigma_0, \ldots, \sigma_l$ of G and any vertex $v \in V(G)$. We define $S(v)$ as the following walk in H:

- If $l = 0$, that is, S is empty, then $S(v) = \varepsilon$.
- If $l = 1$, that is, S contains only one reconfiguration step, then
  $S(v) = \varepsilon$ when $\sigma_0(v) = \sigma_1(v)$ and
  $S(v) = (\sigma_0(v), h)(h, \sigma_1(v))$ otherwise, where h is the color that all neighbors of v have in σ_0 and σ_1.
- If $l > 1$, then $S(v) = S_1(v)S_2(v)\ldots S_l(v)$, where $S_i(v)$ is the walk corresponding to the reconfiguration step from σ_i to σ_{i+1}.

For two H-coloring α, β of G and a vertex $q \in V(G)$, we call an element $Q \in \pi(H)$ *realizable* if there is an H-recoloring sequence $S = \sigma_0, \ldots, \sigma_l$ such that $\sigma_0 = \alpha, \sigma_l = \beta$ and $\overline{S(q)} = Q$. We focus on the following question: which elements of $\pi(H)$ are realizable? It is immediate from the definition that Q must be a reduced walk from $\alpha(q)$ to $\beta(q)$ and have even length (notice that the parity of the length of walks is preserved by reducing, since we only remove pairs of edges ee^{-1}).

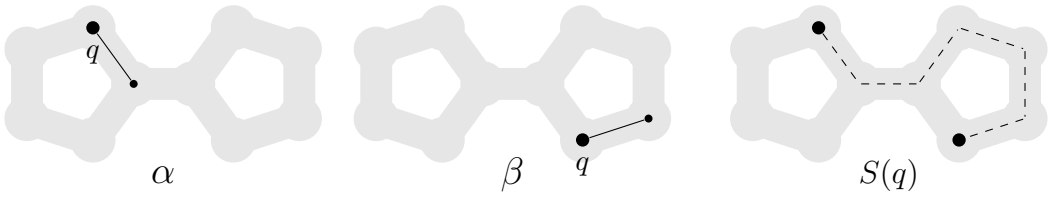

Figure 4: A realizable walk for $\alpha, \beta : K_2 \to H$ and q. Note the shortest walk from $\alpha(q)$ to $\beta(q)$ (of length 3) is not realizable because of parity.

If W is a walk in G and α is an H-coloring of G, then $\alpha(W)$ is a walk in H (one way to see that is to see the walk as a homomorphism from a path: a composition of homomorphisms is a homomorphism).

5. Topological validity

With a homomorphism from G to H one can associate a continuous map from G to H (understood as continuous topological spaces described above). The following lemma is the key to the monochromatic neighborhood property of H. It states in algebraic terms the fact that when an H-coloring is reconfigured into another, the corresponding continuous mappings can be continuously transformed into one another, in a way such that the walk $S(q)$ of any vertex q is the curve traced by the point q.

Lemma 1. *Let $S = \sigma_0, \ldots, \sigma_l$ be an H-recoloring sequence of G. Consider any walk W from vertex u to v in G. Then $\overline{S(v)} = \overline{\sigma_0(W)}^{-1} \cdot \overline{S(u)} \cdot \overline{\sigma_l(W)}$.*

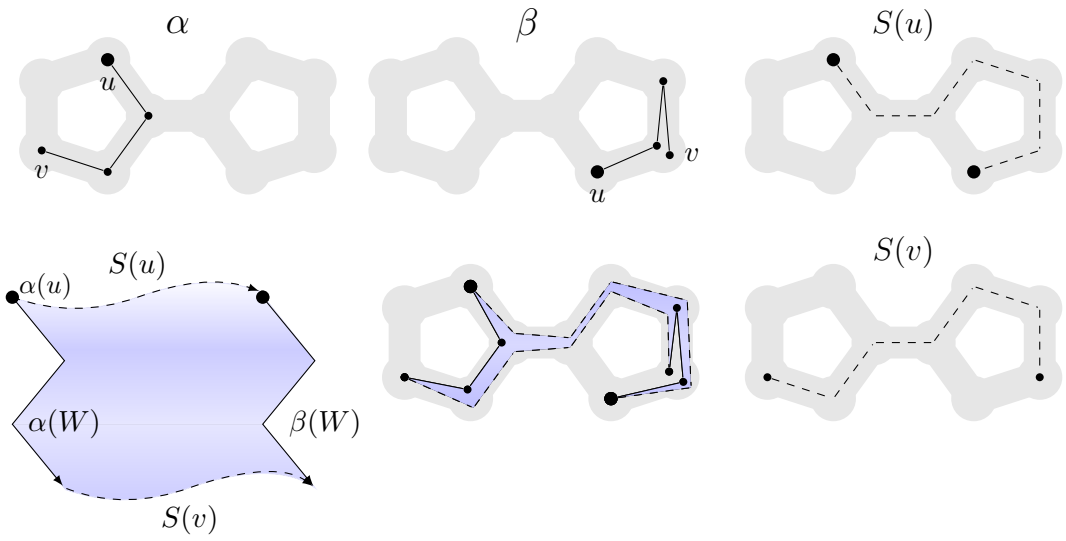

Figure 5: If α can be transformed to β by reconfiguration, then it can by a homotopy $\phi : [0,1] \times [0, |W|]$ such that $\phi(0, \cdot) = \alpha(W)$ and $\phi(1, \cdot) = \beta(W)$. Let $S(u) = \phi(\cdot, 0)$ and $S(v) = \phi(\cdot, |W|)$. Since ϕ is a continuous mapping of a rectangle to H, it's boundary, and so the image of it's boundary, can be contracted to a point: $\overline{\alpha(W)}^{-1} \cdot \overline{S(u)} \cdot \overline{\beta(W)} \cdot \overline{S(v)}^{-1} = \varepsilon$.

We postpone the proof (which follows directly from the monochromatic neighborhood property and induction and is better explained by the 'picture proof' in Figure 5) until the end of the section, to focus on corollaries.

First, we see that in a given instance of H-Recoloring, the reduced vertex walk $\overline{S(q)}$ of one vertex q in a solution S determines all other vertex walks up to reductions (in other words, up to homotopy). Later we will see that in shortest solutions all vertex walks are already reduced, so $\overline{S(q)}$ actually determines the shortest solutions exactly, up to reordering color changes of different vertices. This is why we can focus on the realizability of one element $Q \in \pi(H)$ instead of a whole recoloring sequence.

Second, note that the statement doesn't depend on how we choose walks in G. This means that for every closed walk C in G there is a topological condition on how solutions look like. Intuitively, two solutions can differ only in the number of times they wind around the circuit that C goes around. This is described in more detail in Lemma 15, for now let us only define this condition. The interpretation is that Q is *topologically valid* for α, β, q if there is a homotopy continuously transforming α to β such that q traces the curve Q throughout this transformation ($\phi_0 = \alpha, \phi_1 = \beta$ and the image of $t \mapsto \phi_t(q)$ is Q).

Definition 2. *Let α, β be two H-colorings of G and let q be a vertex of G. Let $Q \in \pi(H)$. We say Q is* topologically valid *for α, β, q if any of the following equivalent statements holds.*

1. *For every closed walk C from q to q we have $\overline{\beta(C)} = Q^{-1} \cdot \overline{\alpha(C)} \cdot Q$.*
2. *For every vertex w and any two walks W_1, W_2 from q to w in G we have $\overline{\alpha(W_1)}^{-1} \cdot Q \cdot \overline{\beta(W_1)} = \overline{\alpha(W_2)}^{-1} \cdot Q \cdot \overline{\beta(W_2)}$.*

Proof (of equivalence). Suppose the second statement is true. Then the first one is true, because C and ε are two walks from q to q, implying

$$\overline{\alpha(C)}^{-1} \cdot Q \cdot \overline{\beta(C)} = \overline{\alpha(\varepsilon)}^{-1} \cdot Q \cdot \overline{\beta(\varepsilon)} = Q.$$

Suppose the first statement is true. To check the second statement let W_1, W_2 be two walks from q to w, for some $w \in V(G)$. Then $W_1 W_2^{-1}$ is a closed walk starting and ending in q, so $\overline{\beta(W_1 W_2^{-1})} = Q^{-1} \cdot \overline{\alpha(W_1 W_2^{-1})} \cdot Q$. Therefore

$$\overline{\alpha(W_2)} \cdot \overline{\alpha(W_1)}^{-1} \cdot Q \cdot \overline{\beta(W_1)} \cdot \overline{\beta(W_2)}^{-1} =$$

$$= \overline{\alpha(W_1 W_2^{-1})}^{-1} \cdot Q \cdot \overline{\beta(W_1 W_2^{-1})} =$$

$$= \overline{\alpha(W_1 W_2^{-1})}^{-1} \cdot Q \cdot Q^{-1} \cdot \overline{\alpha(W_1 W_2^{-1})} \cdot Q =$$

$$= Q.$$

□

Corollary 3. *Let α, β be two H-colorings of G and let q be a vertex of G. If $Q \in \pi(H)$ is realizable for α, β, q then Q is topologically valid for α, β, q.*

It can be shown that this condition also characterizes exactly which curves are realizable as curves drawn by the vertex when continuously transforming α to β, so nothing more can be shown in a purely topological setting.

We note that the existence of a topologically valid walk for α, β implies in particular that cycles have the same homotopy class in α as in β, which algebraically is expressed in the fact that they have the same conjugacy class in the fundamental groupoid (that is, $\overline{\beta(C)} = Q^{-1} \cdot \overline{\alpha(C)} \cdot Q$ for some Q). This generalizes one of the conditions of a characterization of 3-colorings connected by recoloring sequences given by [CHJ11] (Theorem 7 (C2)).

Proof (of Lemma 1). Assume first that $l = 1$, $S = \sigma_0, \sigma_1$, where σ_1 is obtained from σ_0 by recoloring one vertex $w \in V(G)$ from $\sigma_0(w) = a$ to $\sigma_1(w) = b$. Let h be the color that all neighbors of w have in σ_0 and σ_1. By definition $S(w) = (a, h)(h, b)$.

If $W = \varepsilon$ then $u = v$ and the claim follows trivially.

If W has length one, that is $W = (u, v)$, then one of the following holds:

- $u, v \neq w$, implying
  $S(u) = S(v) = \varepsilon$ and $\sigma_0(W) = \sigma_l(W)$.
- $u \neq w$ and $v = w$, implying
  $S(u) = \varepsilon$, $S(v) = (a, h)(h, b)$ and $\sigma_0(W) = (h, a)$, $\sigma_l(W) = (h, b)$.
- $u = w$ and $v \neq w$, implying
  $S(u) = (a, h)(h, b)$, $S(v) = \varepsilon$ and $\sigma_0(W) = (a, h)$, $\sigma_l(W) = (b, h)$.

In either case $\overline{\sigma_l(W)} = \overline{S(u)}^{-1} \cdot \overline{\sigma_0(W)} \cdot \overline{S(v)}$, which is equivalent to the claim.

If W is longer, the claim follows inductively by splitting it into shorter walks: $W = W_1 W_2$ for W_1 from u to some vertex w and W_2 from w to v. It then follows that $\overline{\sigma_l(W)} =$

$$= \overline{\sigma_l(W_1)} \cdot \overline{\sigma_l(W_2)} =$$

$$= \overline{S(u)}^{-1} \cdot \overline{\sigma_0(W_1)} \cdot \overline{S(w)} \cdot \overline{S(w)}^{-1} \cdot \overline{\sigma_0(W_2)} \cdot \overline{S(v)} =$$

$$= \overline{S(u)}^{-1} \cdot \overline{\sigma_0(W_1)} \cdot \overline{\sigma_0(W_2)} \cdot \overline{S(v)} =$$

$$= \overline{S(u)}^{-1} \cdot \overline{\sigma_0(W)} \cdot \overline{S(v)}.$$

It remains to consider the case where S has more than one step. Then the claim follows inductively by writing S as the concatenation of shorter sequences: S_1 which ends in σ_i and S_2 which starts in σ_i. It then follows that $S(v) = S_1(v) S_2(v)$ and $\overline{S(v)} =$

$$= \overline{S_1(v)} \cdot \overline{S_2(v)} =$$

$$= \overline{\sigma_0(W)}^{-1} \cdot \overline{S_1(u)} \cdot \overline{\sigma_i(W)} \cdot \overline{\sigma_i(W)}^{-1} \cdot \overline{S_2(u)} \cdot \overline{\sigma_l(W)} =$$

$$= \overline{\sigma_0(W)}^{-1} \cdot \overline{S_1(u)} \cdot \overline{S_2(u)} \cdot \overline{\sigma_l(W)} =$$

$$= \overline{\sigma_0(W)}^{-1} \cdot \overline{S(u)} \cdot \overline{\sigma_l(W)}.$$

□

6. Tight closed walks and frozen vertices

There is one more obstruction to reconfiguration: closed walks whose sequence of colors are in a sense tightly stretched around H. This generalizes the characterization of frozen vertices in the case of $H = K_3$ from [CHJ11]. In general though, frozen vertices can also arise in other situations, see Figure 6 for an example.

In an H-coloring α of G, a vertex v of G is called *frozen* if for every H-recoloring sequence from α the resulting H-coloring β has $\beta(v) = \alpha(v)$. A closed walk $C = e_1 e_2 \dots e_l$ is *cyclically reduced* if it is reduced and also $e_l \neq e_1^{-1}$. In other words repeating C gives an infinite reduced walk. A closed walk C is α-tight if $\alpha(C)$ is cyclically reduced.

Lemma 4. *Let α be an H-coloring of G and let C be an α-tight walk in G. Then all vertices of C are frozen in α.*

Proof. Suppose to the contrary that there is an H-recoloring sequence $\sigma_0, \dots, \sigma_l$ from α, such that $\sigma_l(C) \neq \sigma_0(C)$. Let i be the least such that $\sigma_i(C) \neq \sigma_0(C)$. Then in σ_{i-1} all vertices of C have the same color as in $\alpha = \sigma_0$ and σ_i is obtained from σ_{i-1} by changing the color of some vertex $v \in C$ from a to b. Let h be the color that all neighbors of v have in σ_{i-1} and σ_i. Let u, w be the vertices of C just before and just after v on C. Since they are neighbors of v, they must have the color h in σ_{i-1} and thus in α. The subsequent edges $(u, v)(v, w)$ of C thus get mapped through α to $(h, a)(a, h)$, contradicting that $\alpha(C)$ is cyclically reduced. □

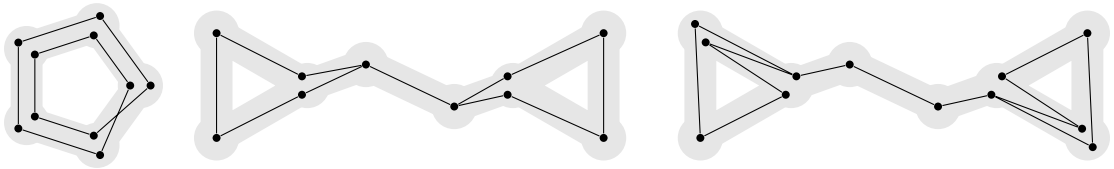

Figure 6: In the left and middle example there is a tight closed walk containing all vertices, no reconfiguration step is possible. In the last example no closed walk is tight, but the 4 middle vertices are frozen.

Tight closed walks can be found by exploring walks W starting from some vertex such that $\alpha(W)$ is reduced. If such walks can be arbitrarily long, or equivalently, contain the same oriented edge twice, then they use an α-tight closed walk. The (potentially infinite) prefix tree of such walks W gives a generalization of the layer construction of [CHJ11]. Algorithmically, we can ask what oriented edges are reachable from any starting oriented edge in time $\mathcal{O}(|E(G)|)$.

Finding a frozen vertex v allows us to limit potentially realizable walks Q to a single one. Since $S(v) = \varepsilon$, Lemma 1 gives $Q = \overline{S(q)} = \overline{\alpha(W)}^{-1} \cdot \overline{S(v)} \cdot \overline{\beta(W)} = \overline{\alpha(W)}^{-1} \cdot \overline{\beta(W)}$.

Corollary 5. *Let $Q \in \pi(H)$ be realizable for α, β – H-colorings of G and $q \in V(G)$. Let C be an $\alpha(C)$-tight closed walk in H. Then for every vertex v on C and every walk W from v to q we have $Q = \overline{\alpha(W)}^{-1} \cdot \overline{\beta(W)}$.*

7. Characterization of realizable walks

In this section we prove that the obstructions described in the previous sections are enough to characterize all possible solutions to an H-Recoloring instance. This is very unexpected, as it shows we can view the graphs as purely topological structures and the only remaining conditions to remember are a simple parity condition and a condition on tight closed walks – the fact that edges are actually discrete and cannot be stretched arbitrarily, for example, turns out to imply no further obstructions to reconfiguration (it only restrict the H-colorings one can be given on the input).

Theorem 6. *Let α, β be two H-colorings of G. Consider any vertex q of G and let $Q \in \pi(H)$ be a reduced walk in H from $\alpha(q)$ to $\beta(q)$. Then Q is realizable for α, β, q if and only if*

- *Q is topologically valid for α, β, q,*
- *Q has even length,*
- *if there is an α-tight walk, then for any vertex v on this walk and any walk W from v to q, $Q = \overline{\alpha(W)}^{-1} \cdot \overline{\beta(W)}$.*

Furthermore, there is a polynomial time algorithm that given G, H, α, β and a walk Q satisfying these conditions, outputs a reconfiguration sequence such that $S(q) = Q$ and $S(v)$ is reduced for all $v \in V(G)$.

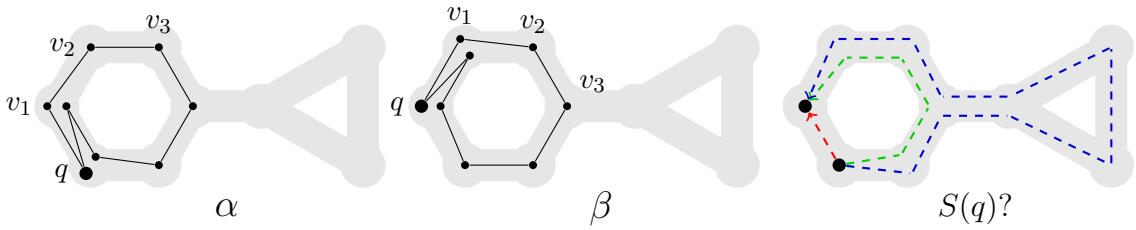

Figure 7: Even though no vertex is frozen, α cannot be reconfigured to β. The red and green walks are not realizable for α, β, q because of parity. The blue walk has good parity, but is not topologically valid (imagine pulling q along it).

Proof. If an H-recoloring sequence is given, then the conditions are satisfied by Corollary 3 and Corollary 5, which proves the 'only if' half.

Consider now a reduced walk $Q \in \pi(H)$ that satisfies the above conditions. For every vertex $v \in V(G)$ let $S_v = \overline{\alpha(W)}^{-1} \cdot Q \cdot \overline{\beta(W)}$ for some walk W from q to v – by topological validity of Q, this doesn't depend on how W is chosen. In particular $S_q = Q$. We will show an H-recoloring sequence S from α to β such that $S(v) = S_v$ for all $v \in V(G)$. The idea is that S_v define a correct H-recoloring sequence for each edge, but it remains to order changes of different vertices into one reconfiguration sequence. Each edge gives a condition on who should recolor first and it is enough to respect these conditions. This is impossible only if there is a cycle of conditions, which turns out to be exactly a α-tight cycle.

Note that since $|Q|$ is even, $|\alpha(W)| = |W| = |\beta(W)|$ and reducing preserves parity, we have that each S_v has even length.

Consider two adjacent vertices $u, v \in V(G)$. Let W be any walk from q to u. Then W followed by (u, v) is a walk from q to v, so $S_v =$

$$= \overline{\alpha(W(u,v))}^{-1} \cdot Q \cdot \overline{\beta(W(u,v))} =$$

$$= (\alpha(v), \alpha(u)) \cdot \overline{\alpha(W)}^{-1} \cdot Q \cdot \overline{\beta(W)} \cdot (\beta(u), \beta(v)) =$$
$$= (\alpha(v), \alpha(u)) \cdot S_u \cdot (\beta(u), \beta(v)).$$

Let

$$S_u = (a_0, a_1)(a_1, a_2)(a_2, a_3) \dots (a_{n-1}, a_n),$$
$$S_v = (b_0, b_1)(b_1, b_2)(b_2, b_3) \dots (b_{m-1}, b_m).$$

Note that using concatenation (as opposed to $\cdot$) we write the exact sequence of edges of a walk. In particular S_u is by definition reduced, so $a_i \neq a_{i+2}$, similarly for S_v. Now $\alpha(u) = a_0, \alpha(v) = b_0$ and $\beta(u) = a_n, \beta(v) = b_m$, so

$$S_v = (b_0, a_0) \cdot S_u \cdot (a_n, b_m)$$

and if we assume S_u is non-empty, then by the parity condition $n \geq 2$ and

$$S_v = \overline{(b_0, a_0)(a_0, a_1)}(a_1, a_2) \dots (a_{n-2}, a_{n-1})\overline{(a_{n-1}, a_n)(a_n, b_m)}.$$

There are two cases, depending on whether $\overline{(b_0, a_0)(a_0, a_1)}$ cancels out to ε. Either it does, that is $(b_0, a_0) = (a_0, a_1)^{-1}$, in which case

$$S_v = (a_1, a_2)(a_2, a_3) \dots (a_{n-2}, a_{n-1})\overline{(a_{n-1}, a_n)(a_n, b_m)}$$

which means $b_0 = a_1, b_1 = a_2, \dots, b_{n-2} = a_{n-1}$ **(case $u \to v$)**
or it doesn't, in which case

$$S_v = (b_0, a_0)(a_0, a_1)(a_1, a_2) \dots (a_{n-2}, a_{n-1})\overline{(a_{n-1}, a_n)(a_n, b_m)}$$

which means $b_1 = a_0, b_2 = a_1, \dots, b_n = a_{n-1}$ **(case $u \leftarrow v$)**.

For two adjacent vertices $u, v \in V(G)$ such that S_u and S_v are non-empty, let us write $u \to v$ in the first case and $u \leftarrow v$ in the other as defined above. We have $u \to v$ iff $v \leftarrow u$, otherwise $b_0 = a_1 = b_2$ (if $u \to v$ and $v \to u$) or $a_0 = b_1 = a_2$ (if $u \leftarrow v$ and $v \leftarrow u$), contradicting that S_u, S_v are reduced walks. We will write S_v^j for the j-th vertex of S_v.

Furthermore, the $\to$ relation has no cycles. Suppose to the contrary that there is a cyclic sequence $v_0, v_1, \ldots, v_{l-1}, v_0$ ($l \geq 3$) such that $v_i \to v_{i+1}$ for $i \in \mathbb{Z}_l$. Then this is an α-tight walk: indeed, arrows imply adjacency in G, and $v_i \to v_{i+1} \to v_{i+2}$ implies that

$$\alpha(v_{i+2}) = S^0_{v_{i+2}} = S^1_{v_{i+1}} = S^2_{v_i} \neq S^0_{v_i} = \alpha(v_i).$$

Therefore by the last condition we have $Q = \overline{\alpha(W)}^{-1} \cdot \overline{\beta(W)}$ and for any walk W from v_i to q,

$$S_{v_i} = \overline{\alpha(W)} \cdot Q \cdot \overline{\beta(W)}^{-1} = \overline{\alpha(W)} \cdot \overline{\alpha(W)}^{-1} \cdot \overline{\beta(W)} \cdot \overline{\beta(W)}^{-1} = \varepsilon.$$

But we didn't assign arrows between vertices whose sequences are empty, a contradiction.

Therefore there is an ordering $v_1, v_2, \ldots, v_{|V(G)|}$ of $V(G)$ such that if $v_i \to v_j$ then $i < j$. We claim the following reconfiguration sequence is a valid H-recoloring sequence from α to β. Recolor:
v_1 from $S^0_{v_1}$ to $S^2_{v_1}$, v_2 from $S^0_{v_2}$ to $S^2_{v_2}$, $\ldots$, v_n from $S^0_{v_n}$ to $S^2_{v_n}$,
v_1 from $S^2_{v_1}$ to $S^4_{v_1}$, v_2 from $S^2_{v_2}$ to $S^4_{v_2}$, $\ldots$, v_n from $S^2_{v_n}$ to $S^4_{v_n}$,
v_1 from $S^4_{v_1}$ to $S^6_{v_1}$, $\ldots$ $\ldots$
We continue this way, disregarding any undefined recolorings to $S^j_{v_i}$ for $j > |S_{v_i}|$. Because of the parity condition, this indeed ends in the coloring β.

To check that it is a valid H-recoloring sequence consider any edge $\{u, v\}$ of G and define a_i, b_i as above. If both S_u and S_v are empty, then $\{u, v\}$ gets constantly mapped to the same edge $\{\alpha(u), \alpha(v)\}$ of H. If exactly one of S_u, S_v is empty, say S_u, then $S_v = (b_0, a_0) \cdot S_u \cdot (a_0, b_m) = (b_0, a_0)(a_0, b_m)$ where $b_0 \neq b_m$ (and $m = 2$). Thus $b_1 = a_0$, so $\{u, v\}$ gets mapped initially to $\{\alpha(u), \alpha(v)\} = \{a_0, b_0\}$ and then to $\{a_0, b_2\} = \{b_1, b_2\}$, which is an edge of H. If both S_u and S_v are non-empty, then assume without loss of generality $u \to v$ (otherwise swap u and v). We have

$$\begin{aligned} S_u &= (a_0, a_1)(a_1, a_2)(a_2, a_3) \ldots (a_{n-2}, a_{n-1})(a_{n-1}, a_n), \\ S_v &= \qquad\quad (a_1, a_2)(a_2, a_3) \ldots (a_{n-2}, a_{n-1})\overline{(a_{n-1}, a_n)(a_n, b_m)} \end{aligned}$$

Thus (u, v) gets mapped initially to $(\alpha(u), \alpha(v)) = (a_0, b_0) = (a_0, a_1)$ and then to $(a_1, a_2), (a_3, a_2), (a_3, a_4), \ldots$, ending in either (a_n, a_{n-1}) or (a_n, a_{n+1}) depending on whether $\overline{(a_{n-1}, a_n)(a_n, b_m)} = \varepsilon$. This is indeed always an edge of H, so since $\{u, v\}$ was chosen arbitrarily, the coloring is always an H-coloring and the sequence is a valid H-recoloring sequence. □

Corollary 7. *Let α, β be two H-colorings of G. Let $S = \sigma_0, \ldots, \sigma_l$ be an H-recoloring sequence between $\sigma_0 = \alpha$ and $\sigma_l = \beta$ such that its length l is minimized. Then for every vertex v, $S(v)$ is reduced.*

Proof. Suppose $S(q)$ is not reduced for some q. Let $Q = \overline{S(q)}$. But the above theorem we know from one side that Q is realizable. From the other side we obtain a solution S' such that $S'(v) = \overline{S'(v)}$ for all v and $\overline{S'(q)} = Q$. By Lemma 1, this implies $S'(v) = \overline{S(v)}$,

for all v. But $\overline{S(v)}$ is always no longer that $S(v)$, and $\overline{S(q)}$ is strictly shorter than $S(q)$. Therefore, since the number of recoloring steps is equal to half the sum of lengths of all $S(v)$, S was not shortest. □

8. An algorithm

The following lemma follows from well-known calculations in the fundamental groupoid, which we recall in Appendix A. Intuitively, we consider for a cycle C of G what this cycle maps to, up to homotopy: in particular $\overline{\alpha(C)}$ must be the same as $\overline{\beta(C)}$ up to homotopy (up to conjugation in $\pi(H)$). Furthermore, if this cycle is non-trivial ($\overline{\alpha(C)} \neq \varepsilon$), then two topologically valid walks may differ only in the number of times they wind around $\overline{\alpha(C)}$ (or the least root R such that $R^n = \overline{\alpha(C)}$). If two non-trivial such cycles wind around different roots, this implies there can be at most one topologically valid walk.

Lemma 8. *Let α, β be H-colorings of G and q a vertex of G. Consider the set $\Pi \subseteq \pi(H)$ of topologically valid walks for α, β, q. One of the following holds:*

0. $\Pi = \emptyset$.

1. $\Pi = \{Q\}$ *for some* $Q \in \pi(H)$.

2. $\Pi = \{R^n \cdot P \mid n \in \mathbb{Z}\}$ *for some* $R, P \in \pi(H)$.

3. Π *contains all reduced walks from* $\alpha(q)$ *to* $\beta(q)$.

Furthermore, there is a algorithm that given α, β, q outputs in polynomial time which case holds and outputs Q or R, P in cases 1,2.

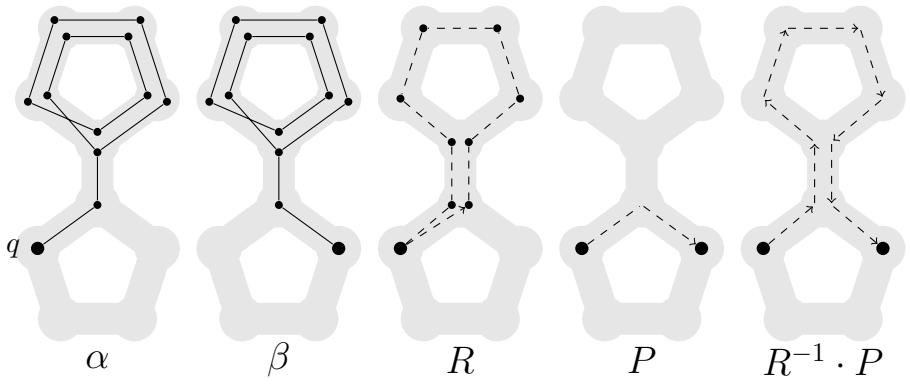

Figure 8: There is essentially only one closed walk C from q to q in this example, which winds around the root R ($\overline{\alpha(C)} = R^2$). The topologically valid paths are exactly $\{R^n \cdot P \mid n \in \mathbb{Z}\}$ (think about deforming α by pulling q: one can pull it once or more around the top cycle by rotating all of α, but this is impossible for the bottom cycle if we are to reach β).

Theorem 9. *Let α, β be H-colorings of G and q a vertex of G. Consider the set $\Pi' \subseteq \pi(H)$ of realizable walks for α, β, q. One of the following holds:*

0. $\Pi' = \emptyset$.

1. $\Pi' = \{Q\}$ *for some* $Q \in \pi(H)$.

2. $\Pi' = \{R^n \cdot P \mid n \in \mathbb{Z}\}$ *for some* $R, P \in \pi(H)$.

3. Π' *contains all reduced walks of even length from* $\alpha(q)$ *to* $\beta(q)$.

Furthermore, there is an algorithm that given α, β, q outputs in polynomial time which case holds and outputs Q or R, P in cases 1,2.

Proof. First find any α-tight closed walk and if there is one, let Q be the only possibly realizable walk as in the last condition of Theorem 6. By running the algorithm from Theorem 6 we can check whether it is realizable and return either $\Pi' = \emptyset$ or $\Pi' = \{Q\}$.

Assume now that there is no α-tight walk. Run the algorithm of Lemma 8 to get a description of topologically valid walks Π and consider the following cases:

0. $\Pi = \emptyset$. Then also $\Pi' = \emptyset$ (see Theorem 6).

1. $\Pi = \{Q\}$ for some $Q \in \pi(H)$. Then return $\Pi' = \{Q\}$ if Q has even length and $\Pi' = \emptyset$ otherwise.

2. $\Pi = \{R^n \cdot P \mid n \in \mathbb{Z}\}$ for some $R, P \in \pi(H)$. The only remaining condition is parity, so one of the following holds:
   - R is even and P is odd – then $\Pi' = \emptyset$,
   - R is even end P is even – then $\Pi' = \Pi$,
   - R is odd and P is even – then $\Pi' = \{R^{2n} \cdot P \mid n \in \mathbb{Z}\}$,
   - R is odd and P is odd – then $\Pi' = \{R^{2n} \cdot (R \cdot P) \mid n \in \mathbb{Z}\}$.

3. Π contains all reduced walks from $\alpha(q)$ to $\beta(q)$. Then Π' contains all reduced walks of even length from $\alpha(q)$ to $\beta(q)$. □

The set of even walks from $\alpha(q)$ to $\beta(q)$ in H is empty if and only if H is bipartite and $\alpha(q), \beta(q)$ are on different sides of a bipartition. This is easily checked in polynomial time, which conludes our main claim. Shortest recoloring sequences can also be found in polynomial time with some more care.

Corollary 10. *Let H be a graph (possibly with loops) with the monochromatic neighborhood property. Then H-Coloring Reachability is in P.*

Theorem 11. *Let H be a graph (possibly with loops) with the monochromatic neighborhood property. Then* Shortest H-Coloring Reachability *is in P.*

Proof sketch. By Corollary 7 all we have to do is minimize

$$\sum_{v \in V(G)} \overline{S(v)} = \sum_{v \in V(G)} |\overline{\alpha(W_v)}^{-1} \cdot Q \cdot \overline{\beta(W_v)}|$$

for $Q \in \Pi'$ from Theorem 9, where W_v is a walk from q to v (arbitrarily chosen). In cases 0,1. this is trivial. In case 2. it is easy to show that $|n| \leq 2|V(G)| + |P|$ in shortest sequences, as repeating R will eventually lengthen all vertex walks.

In case 3. consider a realizable walk Q. Let P_1 be the longest prefix P_1 of Q which will reduce with some $\overline{\alpha(W_v)}^{-1}$ (that is, P_1 is also a prefix of $\overline{\alpha(W_v)}$ for some $v \in V(G)$). Let P_2 bet the longest suffix of Q which will reduce with some $\beta(W_v)$. Either P_1 and P_2 overlap, or $Q = P_1 Q' P_2$ for some $Q' \in \pi(H)$ such that

$$\sum_{v \in V(G)} |\overline{\alpha(W_v)}^{-1} \cdot Q \cdot \overline{\beta(W_v)}| = \sum_{v \in V(G)} \left(|\overline{\alpha(W_v)}^{-1} \cdot P_1| + |Q'| + |P_2 \cdot \overline{\beta(W_v)}| \right).$$

Thus we can guess P_1 by enumerating all prefixes of all $\overline{\alpha(W_v)}$, similarly guess P_2 and guess either an overlap or take Q' to be an arbitrary shortest path of appropriate parity from the tail of P_1 to the head of P_2 in H. □

9. Conclusions and future work

Our result generalizes the algorithm for K_3-Recoloring of [CHJ11] and recovers many of its features in a more general and perhaps more intuitive setting. When limited to $H = K_3$, we see that there is essentially only one possible root R for closed walks in H (a 3-cycle), and the reduced walks in K_3 (between the same endpoints) differ only by powers of R. This can be used to show that either no walk is topologically valid (that is, α, β are not homotopic), or all are. In the second case, it suffices to find tight cycles and either there is one, immediately implying the vertex walks in the only 3-recoloring sequence; or there is none, in which case any vertex walk for one vertex can be realized in a solution, as long as it has even length. In other words, if any walk is realizable, then either $Q = \varepsilon$ is realizable for some (frozen) vertex, or the walk from $\alpha(q)$ to $\beta(q)$ of length 0 or 2 is realizable for any vertex q. In particular, we don't need to find frozen vertices or do any of the calculations in Appendix A, it suffices to run the simple algorithm of Theorem 6 $|V(G)| + 1$ times.

We note that none of the proofs used any structural properties of H. If we consider H-recoloring for any graph H, but only allow recoloring a vertex if all of its neighbors have one common color (in other words, a reconfiguration step is allowed only when the homotopy class of the mapping doesn't change), the same results will follow.

An obvious question is how far can our results be extended to more general CSPs – to the asymmetric case, to multiple constraints, to hypergraphs (relations of arbitrary arity)? Is there any connection with the tractable cases of generalized SAT problems?

Another question is whether the problems of graph homomorphism reconfiguration exhibit a dichotomy. For which graphs H is H-Recoloring in P or PSPACE-complete? For the hard side, it is known that K_4-Recoloring is PSPACE-complete even for G bipartite [BC09]. This is equivalent to saying that H-Recoloring is PSPACE-complete for H the cube graph $K_4 \times K_2$, which similarly implies that H-Recoloring is PSPACE-complete for H the 4-cycle C_4 with all loops added, for example. An easy reduction (known as *folding*, see [FL12]) allows us to focus on so called *stiff* graphs. This is discussed in more detail in the author's master thesis [Wro14a].

Finally, it could be interesting to explore the implications of the monochromatic neighborhood property for the whole Hom complex.

Acknowledgments

The author would like to thank Amer E. Mouawad and Naomi Nishimura for helpful discussions and their hospitality. Many thanks to Jarosław Błasiok for sharing his knowledge of algebraic topology, in a remarkably concise way.

A. Calculations in the fundamental groupoid

In this appendix we prove Lemma 8. Any algorithm will need to limit the number of closed walks considered. The standard way to do that is as follows: fix a vertex $q \in V(G)$ and an arbitrary spanning tree T of G (a minimal connected subgraph that includes all vertices). For $e \in E(G) \setminus E(T)$ and an arbitrary orientation $(\iota(e), \tau(e))$ of e define the *fundamental cycle* C_e as the cycle that goes from q to $\iota(e)$ along the unique shortest path that connects them in T, then to $\tau(e)$ through e, then back to q along the shortest path in T. There are $|E(G)| - |E(T)| = |E(G)| - |V(G)| + 1$ fundamental cycles and together they have the following property (see e.g. Lemma 1.2. of [KN07]).

Fact 12. *Let C be any closed walk from q to q in G. Then $\overline{C} = \overline{C_{e_1}}^{s_1} \cdot \ \ldots \ \cdot \overline{C_{e_n}}^{s_n}$ for some $s_i \in \{-1, +1\}$ and $e_1, \ldots, e_n \in E(G) \setminus E(T)$.*

Fact 13. *There is a polynomial time algorithm that given two H-colorings of a graph G and a vertex q can decide whether there is a $Q \in \pi(H)$ that is topologically valid for α, β, q and if so, outputs such a Q.*

Proof. Q is topologically valid if and only if for every closed walk C from q to q we have $\overline{\beta(C)} = Q^{-1} \cdot \overline{\alpha(C)} \cdot Q$. By Fact 12, this is equivalent to satisfying the equation for each fundamental cycle C_e. Finding a satisfying Q is the Simultaneous Conjugacy Search Problem, for which a linear time algorithm is described in Theorem 6.5. of [MU08]. □

To describe all valid walks we will need the following. For a non-zero closed reduced walk $C \in \pi(H)$ (not necessarily cyclically reduced) we define the *primitive root* of C as the unique $R \in \pi(H)$ such that $C = R^n$ for some $n \in \mathbb{N}$ such that n is maximized. It is a routine exercise to check this is well defined, can be computed in linear time, and that the following holds (see e.g. Lemma 2.1. of [MA80] and Lemma I.2.4. of [Mya04]):

Fact 14. *Let $C_1, C_2 \in \pi(H)$. Then C_1 and C_2 commute, i.e. $C_1 \cdot C_2 = C_2 \cdot C_1$, if and only if $C_1 = \varepsilon$ or $C_2 = \varepsilon$ or both have the same primitive root or one of the roots is the inverse of the other.*

Lemma 15. *Let α, β be two H-colorings of G and let q be a vertex of G. Let $P \in \pi(H)$ be topologically valid for α, β, q Then $Q \in \pi(H)$ is topologically valid for α, β, q if and only if for every closed walk C in G starting and ending in q such that $\overline{\alpha(C)} \neq \varepsilon$, we have*

$$Q = R^n \cdot P$$

for some $n \in \mathbb{Z}$, where $R \in \pi(H)$ is the primitive root of $\overline{\alpha(C)}$.

Proof. Suppose Q is topologically valid and let C be a closed walk starting and ending in q. Then by Lemma 2 we have $\overline{\beta(C)} = P^{-1} \cdot \overline{\alpha(C)} \cdot P$ and $\overline{\beta(C)} = Q^{-1} \cdot \overline{\alpha(C)} \cdot Q$. Therefore

$$P^{-1} \cdot \overline{\alpha(C)} \cdot P = Q^{-1} \cdot \overline{\alpha(C)} \cdot Q$$

$$Q \cdot P^{-1} \cdot \overline{\alpha(C)} = \overline{\alpha(C)} \cdot Q \cdot P^{-1}$$

So $\overline{\alpha(C)}$ commutes with $Q \cdot P^{-1}$. Therefore, if $\overline{\alpha(C)} \neq \varepsilon$ and R is the primitive root of $\overline{\alpha(C)}$, then $Q \cdot P^{-1} = R^n$ for some $n \in \mathbb{Z}$ by Fact 14.

For the other side, suppose that for every closed walk C from q to q with $\overline{\alpha(C)} \neq \varepsilon$ and a primitive root R of $\overline{\alpha(C)}$, there is an $n \in \mathbb{Z}$ such that $Q = R^n \cdot P$. Then for every closed walk C from q to q, $\overline{\alpha(C)}$ commutes with $Q \cdot P^{-1}$, because either $\overline{\alpha(C)} = \varepsilon$ or $\overline{\alpha(C)} = R^k$ and $Q \cdot P^{-1} = R^n$ for some $R \in \pi(H)$ and $n, k \in \mathbb{Z}$. Thus $\overline{\beta(C)} =$

$$= P^{-1} \cdot \overline{\alpha(C)} \cdot P =$$

$$= Q^{-1} \cdot Q \cdot P^{-1} \cdot \overline{\alpha(C)} \cdot P =$$

$$= Q^{-1} \cdot \overline{\alpha(C)} \cdot Q \cdot P^{-1} \cdot P =$$

$$= Q^{-1} \cdot \overline{\alpha(C)} \cdot Q,$$

which shows the topological validity of Q. □

Lemma 8 (restated)**.** *Let α, β be H-colorings of G and q a vertex of G. Consider the set $\Pi \subseteq \pi(H)$ of topologically valid walks for α, β, q. One of the following holds:*

0. $\Pi = \emptyset$.

1. $\Pi = \{Q\}$ *for some* $Q \in \pi(H)$.

2. $\Pi = \{R^n \cdot P \mid n \in \mathbb{Z}\}$ *for some* $R, P \in \pi(H)$.

3. Π *contains all reduced walks from* $\alpha(q)$ *to* $\beta(q)$.

Furthermore, there is a algorithm that given α, β, q outputs in polynomial time which case holds and outputs Q or R, P in cases 1,2.

Proof. Use Fact 13 to compute a topologically valid walk $P \in \pi(H)$ for α, β, q. If there is none, we immediately answer *case 0*. Fix an arbitrary spanning tree and compute the elements $\overline{\alpha(C_e)}$ for all fundamental cycles of G. One of the following holds:

1. There are C_e, C_f such that $\overline{\alpha(C_e)}$ and $\overline{\alpha(C_f)}$ do not commute. Then we show $\Pi = \{P\}$. Clearly $\overline{\alpha(C_i)} \neq \varepsilon$, so let R_i be the primitive root of $\overline{\alpha(C_i)}$ for $i = 1, 2$. Suppose $Q \in \pi(H)$ is topologically valid. Then by Lemma 15, $Q = R_1^{n_1} \cdot P$ and $Q = R_2^{n_2} \cdot P$ for some $n_1, n_2 \in \mathbb{Z}$. Thus $R_1^{n_1} = R_2^{n_2}$. If this element has a primitive root ($n_1, n_2 \neq 0$), then it is equal to both R_1 and R_2, implying that $\overline{\alpha(C_1)}$ and $\overline{\alpha(C_2)}$ have the same primitive root, contradicting Fact 14. Therefore $n_1 = n_2 = 0$ and $Q = P$.
2. There is a C_e such that $\overline{\alpha(C_e)} \neq \varepsilon$ and for every C_f, $\overline{\alpha(C_f)}$ commutes with $\overline{\alpha(C_e)}$. Then by Fact 12, for every closed walk C starting and ending in q, $\overline{\alpha(C)}$ commutes with $\overline{\alpha(C_e)}$. Let R be the primitive root of $\overline{\alpha(C_e)}$. R (or its inverse) is also the primitive root of every non-empty $\overline{\alpha(C)}$, so by Lemma 15, Q is topologically valid iff $Q = R^n \cdot P$ for some $n \in \mathbb{Z}$.
3. For every C_e, $\overline{\alpha(C_e)} = \varepsilon$. Then by Fact 12, for every closed walk C from q to q, $\overline{\alpha(C)} = \varepsilon$. By Lemma 15, every walk $Q \in \pi(H)$ from $\alpha(q)$ to $\beta(q)$ is topologically valid. □

References

[ACORT11] D. Achlioptas, A. Coja-Oghlan, and F. Ricci-Tersenghi. "On the solution-space geometry of random constraint satisfaction problems". *Random Structures and Algorithms* 38.3 (2011), pp. 251–268. arXiv:cs/0611052.

[BC09] P. S. Bonsma and L. Cereceda. "Finding Paths between graph colourings: PSPACE-completeness and superpolynomial distances". *Theoretical Computer Science* 410.50 (2009), pp. 5215–5226.

[BK06] E. Babson and D. N. Kozlov. "Complexes of graph homomorphisms". *Israel Journal of Mathematics* 152.1 (2006), pp. 285–312. arXiv:math/0310056.

[BKW14] P. Bonsma, M. Kaminski, and M. Wrochna. "Reconfiguring Independent Sets in Claw-Free Graphs". *Proceedings of 14th Scandinavian Symposium and Workshops on Algorithm Theory, SWAT 2014*. Vol. 8503. LNCS. 2014, pp. 86–97. arXiv:1403.0359.

[BM14] P. Bonsma and A. E. Mouawad. "The Complexity of Bounded Length Graph Recoloring". *preprint* (2014). arXiv:1404.0337.

[BW04] G. R. Brightwell and P. Winkler. "Graph homomorphisms and long range action". *DIMACS Series in Discrete Mathematics and Theoretical Computer Science* 63 (2004), pp. 29–48.

[CHJ11] L. Cereceda, J. van den Heuvel, and M. Johnson. "Finding paths between 3-colorings". *Journal of Graph Theory* 67.1 (2011), pp. 69–82.

[Dai80] D. P. Dailey. "Uniqueness of colorability and colorability of planar 4-regular graphs are NP-complete". *Discrete Mathematics* 30.3 (1980), pp. 289–293.

[Doc09] A. Dochtermann. "Hom complexes and homotopy theory in the category of graphs". *European J. Combin.* 30.2 (2009), pp. 490–509. arXiv:math/0605275.

[FL12] E. Fieux and J. Lacaze. "Foldings in graphs and relations with simplicial complexes and posets". *Discrete Mathematics* 312.17 (2012), pp. 2639–2651. arXiv:1010.2047.

[FV98] T. Feder and M. Y. Vardi. "The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory". *SIAM Journal on Computing* 28.1 (1998), pp. 57–104.

[Gop+09] P. Gopalan, P. G. Kolaitis, E. N. Maneva, and C. H. Papadimitriou. "The Connectivity of Boolean Satisfiability: Computational and Structural Dichotomies". *SIAM Journal on Computing* 38.6 (2009), pp. 2330–2355. arXiv:cs/0609072.

[HD05] R. A. Hearn and E. D. Demaine. "PSPACE-completeness of sliding-block puzzles and other problems through the nondeterministic constraint logic model of computation". *Theoretical Computer Science* 343.1-2 (2005), pp. 72–96. arXiv:cs/0205005.

[HD09] R. A. Hearn and E. D. Demaine. *Games, puzzles and computation*. A K Peters, 2009.

[HH14] W. Heijltjes and R. Houston. "No Proof Nets for MLL with Units: Proof Equivalence in MLL is PSPACE-complete". *Proceedings of the Joint Meeting of the 23rd EACSL Annual Conference on Computer Science Logic (CSL) and the 29th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. CSL-LICS '14. ACM, 2014, 50:1–50:10.

[HN90] P. Hell and J. Nesetril. "On the complexity of H-coloring". *Journal of Combinatorial Theory, Series B* 48.1 (1990), pp. 92–110.

[Joh+14] M. Johnson, D. Kratsch, S. Kratsch, V. Patel, and D. Paulusma. "Colouring Reconfiguration Is Fixed-Parameter Tractable". *preprint* (2014). arXiv:1403.6347.

[KMM12] M. Kamiński, P. Medvedev, and Martin Milanič. "Complexity of independent set reconfigurability problems". *Theoretical Computer Science* 439 (2012), pp. 9–15. arXiv:1008.4563.

[KN07] J. H. Kwak and R. Nedela. *Graphs and their Coverings*. 2007. URL: http://www.savbb.sk/~nedela/graphcov.pdf.

[MA80] K. Madlener and J. Avenhaus. "String Matching And Algorithmic Problems In Free Groups". *Revista colombiana de matematicas* 14 (1980), pp. 1–16.

[Mou+14] A. E. Mouawad, N. Nishimura, V. Pathak, and V. Raman. "Shortest reconfiguration paths in the solution space of Boolean formulas". *preprint* (2014). arXiv:1404.3801.

[MU08] A. G. Myasnikov and A. Ushakov. "Random subgroups and analysis of the length-based and quotient attacks". *Journal of Mathematical Cryptology* 2.1 (2008), pp. 29–61. arXiv:0707.1501.

[Mya04] A. Myasnikov. *Free groups and elementary equations*. 2004. URL: http://web.cs.gc.cuny.edu/~amyasnikov/Teaching/GT_Spring_2004/Lectures/Lectures.htm.

[Sch13] K. W. Schwerdtfeger. "A Computational Trichotomy for Connectivity of Boolean Satisfiability". *preprint* (2013). arXiv:1312.4524.

[Sch78] T. J. Schaefer. "The complexity of satisfiability problems". *Proceedings of the 10th Annual ACM Symposium on Theory of Computing*. Vol. 78. STOC '78. 1978, pp. 216–226.

[Wro14a] M. Wrochna. "Homomorphism reconfiguration in general graphs". *Reconfiguration and structural graph theory*. Master thesis. 2014. Chap. 4.

[Wro14b] M. Wrochna. "Reconfiguration in bounded bandwidth and treedepth". *preprint* (2014). arXiv:abs/1405.0847.